

BBN Report No. 3359

12

ADA 029368

SPEECH UNDERSTANDING SYSTEMS

Quarterly Technical Progress Report No. 7

1 May 1976 to 31 July 1976

Sponsored by
Advanced Research Projects Agency
ARPA Order No. 2904

DDC
RECEIVED
SEP 4 1976
B

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

This research was supported by the Advanced Research Projects Agency of the Department of Defense and was monitored by ONR under Contract No. N00014-75-C-0533.

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Advanced Research Projects Agency or the U.S. Government.

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER BBN Report No. 3359	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) SPEECH UNDERSTANDING SYSTEMS, Quarterly Technical Progress Report No. 7, 1 May 1976 to 31 July 1976		5. TYPE OF REPORT & PERIOD COVERED Q.T.P.R. 5/1/76 - 7/31/76
7. AUTHOR(s) A. Woods, L. Bates, G. Brown, B. Bruce, C. Cook J. Klovstad, R. Schwartz, J. Wolf		6. PERFORMING ORG. REPORT NUMBER BBN Report No. 3359
9. PERFORMING ORGANIZATION NAME AND ADDRESS Bolt Beranek and Newman Inc. 50 Moulton Street Cambridge, MA 02138		8. CONTRACT OR GRANT NUMBER(s) N00014-75-C-0533, 15 YARPA Order - 6707
11. CONTROLLING OFFICE NAME AND ADDRESS ONR Department of the Navy Arlington, VA 22217		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 5D30 121247.1
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		12. REPORT DATE August 1976
		13. NUMBER OF PAGES 19
		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Distribution of this document is unlimited. It may be released to the Clearinghouse, Department of Commerce for sale to the general public.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) acoustic-phonetic experimental facility, acoustic-phonetics, acoustic-phonetic recognition, acoustics, HWIM, incremental alignment, labeling, lexical retrieval, multi-component systems, natural language retrieval system, natural language understanding, parsing, phoneme discrimination, pragmatic grammar, probabilistic labeling, probabilistic lexical retrieval,		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) This report describes recent progress of the BBN Speech Understanding Systems project covering the period from May 1976 to July 1976. The BBN Speech Understanding project is an effort to develop a continuous speech understanding system which uses syntactic, semantic and pragmatic support from higher level linguistic knowledge sources to compensate for the inherent acoustic indeterminacies in continuous spoken utterances. These knowledge sources are integrated with sophisticated signal processing and		

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

19. Key words (cont'd.)

pronunciation likelihood, question-answering, recognition strategies, resource allocation, segmentation, segmentation errors, segment lattice, semantic interpretation, semantic network, shortfall algorithm, shortfall density, speaker normalization, speaker training, speech, speech data base, speech processing, speech recognition, speech understanding, SUR, synthesis-by-rule, system organization, system performance, word verification.

20. Abstract (cont'd.)

acoustic-phonetic analysis of the input signal, to produce a total system for understanding continuous speech. The system contains components for signal analysis; acoustic parameter extraction; acoustic-phonetic analysis of the signal; phonological expansion of the lexicon; lexical matching and retrieval; syntactic, semantic, and pragmatic analysis and prediction; and inferential fact retrieval and question answering, as well as synthesized text or spoken output.

ACCESSION for	
NTIS	White Section <input checked="" type="checkbox"/>
DOC	Buff Section <input type="checkbox"/>
UNANNOUNCED	<input type="checkbox"/>
JUSTIFICATION	
BY	
DISTRIBUTION AVAILABILITY CODES	
Dist.	Avail. and/or Special
A	

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

SPEECH UNDERSTANDING SYSTEMS
Quarterly Technical Progress Report No. 7
1 May 1976 to 31 July 1976

ARPA Order No. 2904

Contract No. N00014-75-C-0533

Program Code No. 5D30

Principal Investigator:

William A. Woods
(617) 491-1850 x361

Name of Contractor:

Bolt Beranek and Newman Inc.

Scientific Officer:

Marvin Denicoff

Effective Date of Contract:

30 October 1974

Title:

SPEECH UNDERSTANDING SYSTEMS

Contract Expiration Date:

29 October 1976

QTPR Editor:

Bonnie Nash-Webber
(617) 491-1850 x227

Amount of Contract: \$1,966,927

Sponsored by
Advanced Research Projects Agency
ARPA Order No. 2904

This research was supported by the Advanced Research Projects Agency of the Department of Defense and was monitored by ONR under Contract No. N00014-75-C-0533.

Table of Contents

	<u>page</u>
A. Acoustic-Phonetic Recognition	1
B. Lexical Retrieval	3
C. Syntax & Semantics	5
1. Grammar for Semantic Interpretation	5
2. Parser	8
D. Verification	9
E. System & Control Performance	11
1. Recognition Strategies	11
2. System Performance	16
References	18
Appendix	19

A. Acoustic-Phonetic Recognition

Much of this past quarter was spent in designing, implementing, and testing the interface between the Acoustic-Phonetic Experiment Facility (APEF) and the Acoustic-Phonetic Recognition (APR) program. The goal of this interface is to allow the APR to correctly adjust the scores for each phoneme against each segment in the segment lattice according to the particular acoustic feature values found within that segment.

As a test of the effect of this individual adjustment to phoneme scores, we attempted to discriminate among the three nasal consonants. The APR program previously used conventional threshold decisions to choose one of several labels for each segment. The scores for particular phonemes were determined by the statistics of the confusions between these segment labels and the correct phonemes. The highest scoring nasal phoneme was correct 70-75% of the time. With the non-parametric modeling procedure, which uses information from the APEF, the first choice nasal was correct 90% of the time. What is more important is that for those segments which were correct, the scores on the other nasals were often decreased very sharply. For those segments where the first choice was incorrect, the correct nasal had a score near the top scoring nasal.

Speaker Normalization

During the past quarter we also discussed several possible speaker/recording environment normalization procedures. Currently, the APR is speaker independent. That is, any normalizing parameters used are derived directly from the utterance being recognized with no other knowledge about the speaker. While these techniques can perform quite well, the APR could be somewhat more accurate if there were some outside knowledge about this particular speaker. Some of the normalization procedures we discussed are as follows:

- 1) Using a carefully designed, phonetically balanced utterance, one could extract a small number of useful acoustic parameters which were known to be speaker dependent and not easily derivable from an utterance of unknown phonetic content (e.g., average fricative spectra). These could then be used as thresholds in the APR program.

For some uses of a speech understanding system, it would be worth having a trained speech technician extract these numbers. It would also be possible to design a system which would be able to deal with the known utterance, and automatically extract the data.

- 2) Since the statistics of phoneme/segment label confusions in some way reflect the particular speaker characteristics, one could imagine weighting the confusion matrix heavily for that speaker. This would require a large amount of speaker training and would only be useful for some applications. Of course, average statistics could always be used until the speaker's identity or speaker characteristics were determined.
- 3) An extreme case of speaker training would involve deriving the acoustic probability distributions that determine segmentation and labeling from the speech of only one speaker, instead of from a mix of speakers. In principle, the structure of the algorithms

themselves could even vary, though parametric variation with a fixed structure would probably be more practical. Those acoustic recognition programs that use raw spectral matching to determine phonetic content often need this type of speaker tuning [Dixon, 1976, p. 9; Bakis, 1976, p. S97 ; Lowerre, 1976, p. S97].

We have done one experiment on case (2) by using confusion matrix statistics which were heavily weighted toward utterances by the speaker of the utterance, but the results are so far inconclusive. Offsetting the potential advantage to be gained from single-speaker training was a significant reduction in the size of the available set of training utterances and an increased risk of not having important phenomena represented.

B. Lexical Retrieval

During the past quarter, Lexical Retrieval's scoring algorithm was modified to permit a more accurate scoring of alignments that involve segmentation errors. When aligning a phoneme sequence with a segment lattice, the Lexical Retrieval component permits three kinds of "incremental" alignments [Klovstad, 1976]:

- 1) Match - an alignment of one phoneme with one segment.
- 2) Split - an alignment of two consecutive phonemes with one segment.
- 3) Merge - an alignment of one phoneme with two consecutive segments.

Previously, only Matches were scored accurately. During the past quarter, additions were made to three distinct parts of the system in order to permit a more accurate scoring of the Split and Merge alignments.

First we found that in order to calculate Split probabilities, statistics on the frequency of consecutive phonemes were needed. This was accomplished by an extension of the statistics gathering package.

Secondly, we needed to create scoring matrices for both Split and Merge alignments. Let NS be the number of different segments ($NS = 83$ in the current system) and NP , the number of different phonemes ($NP = 105$ in the current system). Then the number of possible Split alignments is $NS \cdot NP \cdot NP$ (915,975) and the number of possible Merge alignments is $NS \cdot NS \cdot NP$ (723,345). Since our data base is somewhat limited and these kinds of segmentation errors occur relatively infrequently (approximately 3 percent of the samples), the possibility of calculating each of these probabilities was clearly out of the question. We wanted to use the alignment statistics available from our data base and restrict ourselves to a more manageable problem. Our solution was to map the segments and phonemes into segment classes and phoneme classes respectively. This permitted the creation of substantially smaller Split and Merge matrices that were indexed on the basis of these classes.

This smeared the statistics somewhat but greatly reduced the size of the matrices. The program that had previously produced the Match scoring matrix was extended to create these two additional matrices.

Thirdly, the search algorithms used by the Lexical Retrieval component had to be modified to use the new probabilistic Split and Merge scores. Other changes such as ones to the alignment programs were also made for system compatibility.

As a result of using these new probabilistic scores, we observed a definite improvement in the overall performance of the Lexical Retrieval component. We expect further improvement as additional sentences become available as part of the data base from which the statistics are gathered since: 1) better estimates of the current "incremental" alignment probabilities will be possible, and 2) mappings to more classes will be possible for the determination of Split and Merge scoring matrices.

C. Syntax & Semantics

1. Grammar for Semantic Interpretation

This quarter the grammar was extensively modified to build semantic interpretations instead of syntactic parse trees. This change was motivated by the fact that with our

grammar encoding semantic as well as syntactic information, producing purely syntactic structures meant semantic information was being thrown away. Such information would then later be reintroduced by the interpretation rules. By eliminating the middle step, we not only speed up the system but reduce its size by an entire TENEX fork.

The grammar now builds interpretations by accumulating in registers the semantic head, quantifier, and links of the nodes being described in the sentence. For example, the sentence

"I will go to Chicago for the ASA meeting."

yields the interpretation

```
(FOR: THE X / (FINDQ: LOCATION (CITY CHICAGO)) : T;
 (FOR: THE Y / (FINDQ: DB/MEETING (SPONSOR ASA)) : T;
  (BUILD: DB/TRIP (DESTINATION X)
    (TRAVELER SPEAKER)
    (TO/ATTEND Y)
    (TIME (AFTER NOW))))
```

This interpretation is built up in the following way. The PUSH arc that looks for a constituent describing a person at the start of the sentence will transform the pronoun "I" into the link-node pair (TRAVELER SPEAKER) and return this as the interpretation of that constituent. The word "will" adds (TIME (AFTER NOW)) to the list of link-node pairs being accumulated. (The grammar does not accept constructions like "will have gone," so "will" can currently always be

interpreted as marking a future event.) The word "go" sets the semantic head to DB/TRIP. The constituent "to Chicago" parses with the interpretation

(DESTINATION (! THE X LOCATION ((CITY CHICAGO)))).

(The ! indicates that a FOR: expression will have to be built as part of the interpretation.) Similarly, "the ASA meeting" produces

(TO/ATTEND (! THE Y DB/MEETING ((SPONSOR ASA)))).

The top level of the grammar has thus accumulated the link-node pairs

(DESTINATION (!---))
(TIME ---)
(TRAVELER SPEAKER)
(TO/ATTEND (!---))

with the semantic head DB/TRIP. The appropriate action (in this case a BUILD:) is created, and the necessary quantificational expressions are expanded around it.

Currently, each level of the grammar produces the regular syntactic parse tree and pops the semantic interpretations that it has built in parallel as a feature. As soon as the grammar has been thoroughly checked out, the syntactic registers and parse trees will be eliminated.

This major change to the grammar has necessitated a number of changes to the semantic network. For example, the per diem associated with a city had been represented as a property (one-way link) of the city. In English, however, the per diem is usually referenced by a noun phrase as in "What is the per diem for Chicago?" The most natural interpretation results in a per diem structure, which itself has properties (e.g., a dollar value) and relations (e.g., an associated city). We are now restructuring the relevant parts of the network to be compatible with the resultant interpretations. At the same time the network is being enlarged to include all the place and people names in BIGDICT. There are now approximately 2400 nodes in the network.

In addition to network changes, there are also a number of changes being made to the retrieval functions and new METHODS are being added for the "fictitious links" that appear in interpretations [Bruce and Harris, 1975].

2. Parser

During the past quarter, work on the parser has centered on fixing bugs, implementing a facility for handling island collision events, and designing modifications to increase the number of syntactic events that can be processed (including a garbage collector for

path configurations and a facility for swapping some of the data arrays). Features have also been added to provide for the computation of a syntactic likelihood score by actions on the arcs of the grammar. This syntactic likelihood score will be used initially to provide score adjustments to events as a result of confirmation or disconfirmation of prosodic hypotheses made by arcs of the grammar. The implementation of this facility is one of the steps necessary to the incorporation of the UNIVAC boundary detection programs into the system, a step that we hope to be able to try.

In addition, the parser has been modified to permit the lifting of registers from one level to the next as features in order to pass along semantic interpretations.

D. Verification

During the past quarter, we extended the scoring mechanism of the Verification component to provide log-likelihood ratios of verified word scores. To do this, spectral distance scores (old scores) were collected from 300 words that had been verified by the speech understanding system. Of these 300 words, approximately 50 were correct, "correct" being defined as having verified the proper phonetic spelling over the appropriate region of the utterance. We created two histograms based on old scores,

one for correct words only and the other for all words. We then modeled each of the the histograms according to [Makhoul and Schwartz, 1975, pp. 50-65] in order to provide smooth continuous probability density functions. These two models were entered into the Verification component. We now compute the old score as before, then divide the probability of that score for correct words by the probability of that score for all words. We take the log of this ratio which gives us the log-likelihood ratio or new score. The control component has been appropriately modified to accept this score and combine it with the log-likelihood score returned by the Lexical Retrieval component.

A new synthesis-by-rule program has been developed during this quarter, differing from its predecessor in that it produces synthesis parameters to drive a linear predictive waveform synthesizer. This was done in order to make the synthesis output more compatible with the error metric used in the verification component. At the quarter's end, a new version of the Verification component based on this synthesis program was being assembled.

E. System and Control Performance

1. Recognition Strategies

During the past quarter, we continued our development of shortfall control strategies and built experimental versions of the system corresponding to each of three of them. In addition, we implemented an initial version of a new control strategy, called a "bounded-breadth left-end" strategy, resulting in a fourth version of the system. Performance results for these four evolutionary stages of the system are given in the next section.

The first three systems differ from previous ones primarily in their control strategy. (These changes will be described briefly below.) The Acoustic-Phonetic Recognizer, Lexical Matcher, and Dictionary (i.e., phonological rules) are basically unchanged, although part way through the June 18 system testing, it was discovered that the APR confusion statistics in use had been computed using an algorithm that was thought to have been rejected some time ago. A change to the "correct" confusion statistics was made instantly. With respect to the higher-level components, the only differences lie in small changes made to the details of the grammar.

In the first of these three systems, referred to as the "May 23 system," the "credit" heuristic [Woods, 1976, pp. 138-141] was added to the shortfall density method of computing priority scores on events. In addition, the score of a word-match was changed to take into account the "pronunciation likelihood" score of the given pronunciation. This is derived during the dictionary expansion phase from likelihoods of phonological rules being applied during the expansion process.

In the second of these systems, called the "June 6 system," island collision events were added. That is, as words are added to a theory, checks are performed to see if the added words correspond to words previously added to other theories in the opposite direction. For each such "collision," an event is made that merges the two events, and it takes its place on the event queue with a score appropriate to the word matches in the combined hypothesis. As with all events, the syntactic consistency of the new event is not checked unless and until the event becomes the top element on the event queue.

Another significant change was in the "rectification" of adjacent word matches in a theory -- i.e., the rejection of paths that use incompatible adjacent word matches. Formerly, the score of a series of adjacent (fuzzy) word matches was the sum of the best individual word match in

each fuzzy, regardless of whether or not they were consistent with each other. In the June 6 system, only word matches with common boundaries are allowed to abut, and the scoring of adjacent fuzzy word matches involves examining the allowable word match pairs and picking the consistent path with the best score. Also, priority scoring was computed using shortfall density alone, with neither credit nor liability.

In the June 18 system, this concept of rectification was extended to involve not just boundary consistency, but full phonological consistency. That is, control is now aware of which member of a context fuzzy word match has been an anchor for each new word match, and it uses this information in rectification. This brings to bear the full effect of the word-boundary phonological rules employed in the Lexical Matcher. In addition, several bugs were fixed in the computation of shortfall and credit scores, and priority scores are once again computed as shortfall density plus credit. However, a major bug remains in the June 18 system affecting the scoring of seed words with phonological word boundary effects. This should be fixed in a later version.

One of the problems with the shortfall control strategies is the large number of events that must usually be processed in order for one theory to grow large enough to

span the entire utterance. Several seeds are effectively started in parallel, and their successive generations of descendents grow rapidly. The usual mode of failure of our shortfall systems was for the system to run out of space in either the control or syntax fork, or for it to hit an arbitrarily imposed 2 cpu-hour time-out. By this time, the system would usually have processed between 60 and 100 events.

For these reasons, we have also implemented a different type of control strategy, which we have dubbed a "bounded-breadth left-end" strategy. In essence, the procedure is as follows:

1. Scan for possible utterance-initial words at the left end of the utterance. Form an initial event queue from the resulting seed events.
2. Order the event queue by event score (word match quality), then discard all but the best N events.
3. If all events span the utterance, go to step 3b.
 - 3a. Select for syntactic processing the event whose duration is the shortest, but do not select any event that spans the utterance.
 - 3b. Select the top (best scoring) event.
4. Give that event to Syntax for syntactic verification.
 - 4a. If the event spans and is linguistically well-formed and complete, declare that to be the interpretation.
 - 4b. If the event is rejected as ill-formed, go to step 2.

4c. If Syntax proposes words and classes that might occur to the right, give the proposals to Lexical Retrieval. Form new events from any word matches that come back and add them to the event queue. Go to step 2.

This strategy effectively amounts to starting off N events at the left end of the utterance and forcing the best N (or fewer) events at each point all the way to the right end. Once all events hit a possible right end boundary, the highest scoring syntactically acceptable event (if any) is declared the winner.

This remarkably simple control strategy has an upper bound on the number of events to be processed, on the order of N times the number of words in the utterance. This is clearly linear with respect to the length of the utterance, not exponential, as in the island-driven strategies. Also, the partial interpretations of the utterance are anchored to the left end, which provides rather tighter syntactic constraints at each step than is the case with the middle-out strategies. The disadvantages, of course, are that the system must find the leftmost word in the initial scan, and it has no more than one chance at each choice point to find each successive correct word, and furthermore, to find it with a sufficiently good score for it to be kept on the queue of maximum length N .

Such a "bounded-breadth left-end" control strategy was made an option in the "July 22 system." Given the way in which HWIM's control component is implemented, the inclusion of this strategy required the addition of only a handful of functions and the setting of several existing option flags. With the maximum queue length N set to 8, and utterances of 3 to 9 words each, the system only rarely runs out of resources before terminating. The July 22 system also included addition to the Lexical Retrieval component of probabilistic split and merge scoring, as described in Section B.

2. System Performance

We have been testing system performance on several sets of utterances: (a) the three sets of 20 utterances each selected by SCRL, and designated by us as the "March", "April", and "May", sets; (b) six of the "May" utterances, re-recorded in a very quiet room, designated the "June" set. These were re-recorded to test the hypothesis that the higher noise level in our new laboratory is detrimental to the operation of the acoustic-phonetic recognition; (c) We also have been using two other sets of 10 utterances each, from our collection of on-line utterances dating before February 1976. Since some of these utterances were used to tune the APR, we do not regard results obtained with them as being indicative of system performance on new utterances.

"Control set #1" is made up of 10 utterances on which the March-vintage systems succeeded; "control set #2" is made up of utterances on which the March-vintage systems failed. So the first control set represents "good" utterances on which we should expect to continue to succeed; the second represents utterances on which new successes are sought.

The utterance-successes for the four versions of the system are summarized below, where the four dates heading the columns represent the versions of the system described in the previous section.

	<u>May 23</u>	<u>June 6</u>	<u>June 18</u>	<u>July 22</u>
March	--	--	3/20	7/20
April	--	--	2/20	4/20
May	1/20	1/20	1/20	5/20
June	--	--	1/6	0/6
TOTAL (M-J)			7/66=11%	16/66=24%
C.S.#1	8/10	--	7/10	8/10
C.S.#2	0/10	--	0/10	1/10

References

Bakis, R. (1976)

"Continuous Speech Recognition via Centisecond Acoustic States," (Computer Science Dept., IBM Watson Research Center) presented at the 91st Meeting of ASA, April, 1976.

Bruce, B. and G. Harris (1975)

"Procedural Semantics in the Travel System," in Speech Understanding Systems, Quarterly Technical Progress Report No. 3, 1 May to 31 July 1975, Report No. 3115, Bolt Beranek and Newman Inc., Cambridge, MA 02138.

Dixon, N.R. and H.F. Silverman (1976)

"The 1976 Modular Acoustic Processor (MAP): Signal Analysis and Phonemic Segmentation," IEEE - ICASSP, April, 1976

Klovstad, J. (1976)

"Probabilistic Lexical Retrieval with Embedded Phonological Word Boundary Rules," in Speech Understanding Systems, Quarterly Technical Progress Report No. 6, 1 February to 30 April 1976, Report No. 3303, Bolt Beranek and Newman Inc., Cambridge, MA, 02138.

Lowerre, P. and B. Raj Reddy (1976)

"Harpy, a Connected Speech Recognition System," (Dept. of Computer Science, CMU) presented at the 91st Meeting of the ASA, April, 1976

Makhoul, J. and R. Schwartz (1975)

"Modeling of Probability Distributions," in Speech Understanding Systems, Annual Technical Report, 30 October 1974 to 29 October 1975, Report No. 3188, Bolt Beranek and Newman Inc., Cambridge, MA, 02138.

Woods, W.A. (1976)

"Shortfall Scoring Strategies for Speech Understanding Control," in Speech Understanding Systems, Quarterly Technical Progress Report No. 6 (see above).

APPENDIX

Publications

Wolf, J.J., "Knowledge Hypotheses and Control in the HWIM Speech Understanding System," Joint Conference on Pattern Recognition and Artificial Intelligence, Hyannis, MA, 1-3 June 1976.

Abstract

The BBN Speech Understanding System (dubbed HWIM, or Hear What I Mean) contains knowledge sources at the levels of acoustic-phonetics, phonology, vocabulary, syntax, semantics, factual knowledge, and discourse. This paper describes how these knowledge sources are realized in the nine functional components of the system. It also discusses the control strategy for generating, evaluating, and extending hypotheses into a complete understanding of the spoken utterance.

Nash-Webber, B.L., "Semantic Interpretation Revisited," BBN Report No. 3335, Bolt Beranek and Newman Inc., Cambridge, MA. Also, presented at COLING-76, Ottawa, Canada, 28 June-2 July 1976.

Abstract

A brief overview is given of the BBN LUNAR system. This is followed by a discussion of two of its deficiencies: the simple "enumerate and test" processing of quantified expressions in its meaning representation language and its inadequate treatment of anaphora. We then present a rough classification of anaphoric expressions as groundwork towards formulating a general computational treatment of the phenomenon. In this classification, we establish a distinction between denotational anaphora - references to previously mentioned objects, sets, events, states, etc. - and descriptive anaphora - references to previous descriptions. Finally, we present an initial sketch of both a formal meaning representation language and some procedures seen needed for manipulating sentences of that language, which may provide a handle on some aspects of anaphor resolution.

Nash-Webber, B.L. and B. Bruce, "Evolving Uses of Knowledge in a Speech Understanding System," presented at COLING-76, Ottawa, Canada, 28 June-2 July 1976.

Bruce B., "Discourse Influences on Language Generation," presented at COLING-76, Ottawa, Canada, 28 June-2 July 1976.

Official Distribution List
Contract N00014-75-C-0533

Defense Documentation Center
Cameron Station
Alexandria, VA 22314

Office of Naval Research
Information Systems Program
Code 437
Arlington, VA 22217

Office of Naval Research
Code 1021P
Arlington, VA 22217

Office of Naval Research
Branch Office, Boston
495 Summer Street
Boston, MA 02210

Office of Naval Research
Branch Office, Chicago
536 South Clark Street
Chicago, ILL 60605

Office of Naval Research
Branch Office, Pasadena
1030 East Green Street
Pasadena, CA 91106

New York Area Office
715 Broadway - 5th floor
New York, N.Y. 10003

Naval Research Laboratory
Technical Information Division
Code 2627
Washington, D.C. 20375

Dr. A. L. Slafkosky
Scientific Advisor
Commandant of the Marine Corps
Code RD-1
Washington, D.C. 20380

Office of Naval Research
Code 455
Arlington, VA 22217

Office of Naval Research
Code 458
Arlington, VA 22217

cont'd.

Naval Electronics Laboratory Center
Advanced Software Technology Division
Code 5200
San Diego, CA 92152

Mr. E. H. Gleissner
Naval Ship Research & Development Center
Computation and Mathematics Department
Bethesda, MD 20084

Captain Grace M. Hopper
NAICOM/MIS Planning Branch (OP-916D)
Office of Chief of Naval Operations
Washington, D.C. 20350

Mr. Kin B. Thompson
Technical Director
Information Systems Division (OP-91T)
Office of Chief of Naval Operations
Washington, D.C. 20350

Advanced Research Projects Agency
Information Processing Techniques
1400 Wilson Boulevard
Arlington, VA 22209

Commanding Officer
Naval Air Development Center
Warminster, PA 18974

Professor Omar Wing
Dept. of Electrical Engineering
& Computer Science
Columbia University in the City of New York
New York, N.Y. 10027

Assistant Chief for Technology
Office of Naval Research
Code 200
Arlington, VA 22217